arXiv:1402.5398v1 [math.NA] 21 Feb 2014

An exact solver for simple $\mathcal{H}$-matrix systems

Steffen Börm and Jessica Gördes

February 24, 2014

Hierarchical matrices (usually abbreviated $\mathcal{H}$-matrices) are frequently used to construct preconditioners for systems of linear equations. Since it is possible to compute approximate inverses or LU factorizations in $\mathcal{H}$-matrix representation using only $\mathcal{O}(n \log^2 n)$ operations, these preconditioners can be very efficient.

Here we consider an algorithm that allows us to solve a linear system of equations given in a simple $\mathcal{H}$-matrix format *exactly* using $\mathcal{O}(n \log^2 n)$ operations. The central idea of our approach is to avoid computing the inverse and instead use an efficient representation of the LU factorization based on low-rank updates performed with the well-known Sherman-Morrison-Woodbury equation.

1 Introduction

Hierarchical matrices have been introduced in [6, 7] as a technique for representing certain dense matrices in a data-sparse and therefore efficient way. The approach is related to the well-known multipole [11, 5] and panel clustering [8, 9] techniques: instead of approximating a smooth function by a degenerate expansion, a matrix block is approximated by a low-rank matrix. The algebraic approach offers the possibility to perform matrix arithmetic operations efficiently and to treat general matrices.

Already the first papers on $\mathcal{H}$-matrices, e.g., [6], consider the question of solving linear systems of equations with a system matrix given in $\mathcal{H}$-matrix form. Until now, the standard approach has been to compute an approximation of the inverse [4] or at least an approximate LU factorization [3, 1]. Combined with a well-chosen clustering strategy, particularly the LU factorization can be very efficient and rivals algebraic multigrid algorithms [10].

Still, even the most refined LU factorization is based on the $\mathcal{H}$-matrix multiplication algorithm, and this algorithm typically finds only approximations, although these approximations can be arbitrarily accurate.

In this paper, we present an algorithm that solves a system of linear equations given in a simple $\mathcal{H}$-matrix representation *exactly*, at least up to rounding errors introduced by floating point arithmetic operations. The algorithm is based on the LU factorization,

but while standard algorithms form the Schur complement explicitly, we handle it implicitly using the Sherman-Morrison-Woodbury formula. Due to this approach, the local ranks are preserved, no truncation to lower rank is required, and therefore the resulting decomposition can be used to solve the system directly.

The algorithm can be split into two phases: a setup step computes the quantities describing the factorization of the matrix, and a solver step then solves the linear system. The first step requires $\mathcal{O}(n \log^2 n)$ operations, where n is the matrix dimension, and has to be carried out only once for a given matrix. The second step requires only $\mathcal{O}(n \log n)$ operations and computes the solution for a given right-hand side.

It should be mentioned that there are other algorithms for solving similar problems: if the matrix is hierarchically semi-separable, it is possible to solve systems in $\mathcal{O}(n)$ operations [2], but this works only if the low-rank blocks are of a very special nested structure, not for more general $\mathcal{H}$-matrices.

2 Matrix structure and basic idea

In order to keep the presentation of the basic ideas simple, we restrict our attention to the simplest $\mathcal{H}$-matrix structure [6]:

Definition 2.1 ($\mathcal{H}$-matrix) *Let $n_0 \in \mathbb{N}$. We let*

$$\mathcal{H}_0 := \mathbb{R}^{n_0 \times n_0}$$

and define $\mathcal{H}$-matrices on higher levels inductively: let $\ell \in \mathbb{N}$ and $n_\ell := n_0 2^\ell$. A matrix $A \in \mathbb{R}^{n_\ell \times n_\ell}$ is an element of $\mathcal{H}_\ell \subseteq \mathbb{R}^{n_\ell \times n_\ell}$ if and only if there are matrices $A_1, A_2 \in \mathcal{H}_{\ell-1}$ and vectors $a_1, a_2, b_1, b_2 \in \mathbb{R}^{n_{\ell-1}}$ satisfying

$$A = \begin{pmatrix} A_1 & a_1 b_1^* \\ a_2 b_2^* & A_2 \end{pmatrix}. \tag{1}$$

We call the set $\mathcal{H}_\ell$ the set of $\mathcal{H}$-matrices on level ℓ.

Given an $\mathcal{H}$-matrix $A \in \mathcal{H}_\ell$ and a right-hand side vector $z \in \mathbb{R}^{n_\ell}$, we are interested in finding $x \in \mathbb{R}^{n_\ell}$ with

$$Ax = z. \tag{2}$$

In general, this is only possible if A is regular. Since our algorithm uses a hierarchy of sub-problems to solve the system, we require A to have a more restrictive property:

Definition 2.2 (Hierarchically regular) *Let $A \in \mathcal{H}_\ell$. We call A* hierarchically regular *if it is regular and, in case $\ell > 0$, if the submatrices $A_1, A_2 \in \mathcal{H}_{\ell-1}$ of its representation (1) are also hierarchically regular.*

We can see that, e.g., positive definite matrices are hierarchically regular, since all of their diagonal blocks are positive definite and therefore regular.

Let $\ell \in \mathbb{N}$. If $A \in \mathcal{H}_\ell$ is hierarchically regular, its block LU decomposition is given by

$$A = \begin{pmatrix} A_1 & a_1 b_1^* \\ b_2 a_2^* & A_2 \end{pmatrix} = \begin{pmatrix} I & \\ b_2 a_2^* A_1^{-1} & I \end{pmatrix} \begin{pmatrix} A_1 & a_1 b_1^* \\ & A_2 - b_2 a_2^* A_1^{-1} a_1 b_1^* \end{pmatrix},$$

and we can use this decomposition to solve the linear system (2). Note that, since A is hierarchically regular, the matrices A and A_1 are regular, therefore the Schur complement $A_2 - b_2 a_2^* A_1^{-1} a_1 b_1^*$ also has to be regular.

In order to make handling the Schur complement easier, we introduce

$$c_A := (A_1^{-1})^* a_2, \qquad \gamma_A := c_A^* a_1 = a_2^* A_1^{-1} a_1 \tag{3}$$

and get

$$A = \underbrace{\begin{pmatrix} I & \\ b_2 c_A^* & I \end{pmatrix}}_{=:L} \underbrace{\begin{pmatrix} A_1 & a_1 b_1^* \\ & A_2 - \gamma_A b_2 b_1^* \end{pmatrix}}_{=:U} = LU.$$

Now we can consider solving the linear system by block forward and backward substitution, i.e., we will solve

$$Ly = z, \qquad Ux = y.$$

We split the vectors x, y and z into subvectors $x_1, x_2, y_1, y_2, z_1, z_2 \in \mathbb{R}^{n_{\ell-1}}$ with

$$x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \qquad y = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}, \qquad z = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}, \tag{4}$$

and can write $Ly = z$ in the form

$$\begin{pmatrix} I & \\ b_2 c_A^* & I \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}, \qquad y_1 = z_1, \qquad y_2 = z_2 - b_2 c_A^* y_1.$$

Solving $Ux = y$ for x is a little more involved, since we have

$$\begin{pmatrix} A_1 & a_1 b_1^* \\ & A_2 - \gamma_A b_2 b_1^* \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}, \quad (A_2 - \gamma_A b_2 b_1^*) x_2 = y_2, \quad A_1 x_1 = y_1 - a_1 b_1^* x_2$$

and have to find a way of solving both sub-problems efficiently.

In order to handle the first equation, we rely on the well-known Sherman-Morrison-Woodbury equation [12]. In our case, it yields

$$\left(I + \frac{\gamma_A A_2^{-1} b_2 b_1^*}{1 - \gamma_A b_1^* A_2^{-1} b_2} \right) A_2^{-1} = (A_2 - \gamma_A b_2 b_1^*)^{-1}.$$

We simplify the equation by introducing

$$d_A := A_2^{-1} b_2, \qquad \delta_A := \gamma_A b_1^* d_A = \gamma_A b_1^* A_2^{-1} b_2 \tag{5}$$

and get

$$\left(I + \gamma_A \frac{d_A b_1^*}{1 - \delta_A}\right) A_2^{-1} = (A_2 - \gamma_A b_2 b_1^*)^{-1}, \tag{6}$$

and x_2 can be computed by first recursively finding $\widehat{x}_2 \in \mathbb{R}^{n_{\ell-1}}$ with

$$A_2 \widehat{x}_2 = y_2$$

and then using the rank one correction

$$x_2 = \widehat{x}_2 + \gamma_A \frac{b_1^* \widehat{x}_2}{1 - \delta_A} d_A.$$

Once x_2 has been computed, we can proceed to recursively solve

$$A_1 x_1 = y_1 - a_1 b_1^* x_2$$

to determine x_1, and therefore the solution x.

Of course we also need efficient algorithms for computing the auxiliary vectors c_A and d_A introduced in (3) and (5). Since (3) involves the inverse of the adjoint of A_1, we require an algorithm for solving systems of the form

$$A^* x = z. \tag{7}$$

Fortunately, we can use the LU factorization to solve this problem as well: due to $A = LU$, we also have $A^* = U^* L^*$ and can solve

$$U^* y = z, \qquad\qquad L^* x = y$$

by forward and backward substitution. Using the subvectors defined in (4), the forward substitution takes the form

$$\begin{pmatrix} A_1^* & \\ b_1 a_1^* & A_2^* - \bar{\gamma}_A b_1 b_2^* \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}, \quad A_1^* y_1 = z_1, \quad (A_2^* - \bar{\gamma}_A b_1 b_2^*) y_2 = z_2 - b_1 a_1^* y_1.$$

We can compute y_1 by recursion and use the adjoint of equation (6) to get

$$(A_2^*)^{-1} \left(I + \bar{\gamma}_A \frac{b_1 d_A^*}{1 - \bar{\delta}_A}\right) = (A_2^* - \bar{\gamma}_A b_1 b_2^*)^{-1},$$

and this allows us to compute y_2 in the form

$$\widehat{z}_2 = z_2 - b_1 a_1^* y_1, \qquad \widehat{y}_2 = \widehat{z}_2 + \bar{\gamma}_A \frac{d_A^* \widehat{z}_2}{1 - \bar{\delta}_A} b_1, \qquad A_2^* y_2 = \widehat{y}_2.$$

Now we can turn our attention to the backward substitution to solve

$$\begin{pmatrix} I & c_A b_2^* \\ & I \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}, \qquad x_2 = y_2, \qquad x_1 = y_1 - c_A b_2^* x_2,$$

which fortunately requires only inner products and linear combinations.

3 Algorithm and complexity

We have seen that we can compute the solution of the systems (2) and (7) efficiently if we are able to solve sub-problems involving the two diagonal blocks A_1 and A_2 and their adjoints. Assuming that the auxiliary vectors c_A and d_A and the values γ_A and δ_A have already been prepared, this leads to the algorithm given in Figure 1.

procedure solve(A, **var** x);
begin
  if $\ell = 0$ **then**
    Solve directly
  else begin
    $\alpha_1 \leftarrow c_A^* x_1$; $x_2 \leftarrow x_2 - \alpha_1 b_2$
    solve(A_2, x_2)
    $\alpha_2 \leftarrow b_1^* x_2$; $\alpha_3 \leftarrow \gamma_A \alpha_2 / (1 - \delta_A)$; $x_2 \leftarrow x_2 + \alpha_3 d_A$
    $\alpha_4 \leftarrow b_1^* x_2$; $x_1 \leftarrow x_1 - \alpha_4 a_1$
    solve(A_1, x_1)
  end
end

Figure 1: Solve the linear system: On entry, the vector x contains the right-hand side of problem (2). Recursive solves and low-rank updates are used to replace it by the solution. We assume that the vectors c_A and d_A and the values γ_A and δ_A have already been prepared.

The algorithm is called with $x = z$ and overwrites the vector x with the solution of system (2). If $A \in \mathcal{H}_0$, the matrix can be considered small and we can solve the system directly. If $A \in \mathcal{H}_\ell$ for $\ell > 0$, the recursive procedure described in the previous section is used: the first line corresponds to the forward substitution in L and overwrites x_2 by y_2. In the second line, we recursively solve a linear system with the matrix A_2 to overwrite x_2 by $\widehat{x}_2$. In the third line, we perform the Sherman-Morrison-Woodbury update to get the "lower" half x_2 of the solution vector. In the fourth and fifth line, its "upper" half x_1 is computed by first updating the right-hand side and then recursively solving the remaining system.

The adjoint system (7) can be solved in a similar fashion by using $A^* = U^* L^*$ as described in the previous section, this leads to the algorithm given in Figure 2.

Both algorithms work only if the auxiliary vectors c_A and d_A and the auxiliary values γ_A and δ_A have already been prepared. Fortunately, computing c_A for a matrix $A \in \mathcal{H}_\ell$ requires only solving the adjoint system for $A_1^* \in \mathcal{H}_{\ell-1}$, and similarly d_A can be computed by solving the system for $A_2 \in \mathcal{H}_{\ell-1}$. This means that we can prepare these vectors by bootstrapping: on level $\ell = 0$, we do not require the vectors, but we may want to prepare auxiliary structures for solving efficiently, e.g., by computing a suitable factorization of the matrix A. On level $\ell = 1$, we have to solve systems on level $\ell - 1 = 0$ in order to find c_A and d_A, but this can be done directly. Once the vectors on a level ℓ have been

```
procedure solveadj(A, var x);
begin
  if ℓ = 0 then
    Solve directly
  else begin
    solveadj(A_1, x_1)
    α_1 ← a_1^* x_1;   x_2 ← x_2 − α_1 b_1
    α_2 ← d_A^* x_2;   α_3 ← γ̄_A α_2/(1 − δ̄_A);   x_2 ← x_2 − α_3 b_1
    solveadj(A_2, x_2)
    α_4 ← b_2^* x_2;   x_1 ← x_1 − α_4 c_A
  end
end
```

Figure 2: Solve the linear system: On entry, the vector x contains the right-hand side of adjoint problem (7). Recursive solves and low-rank updates are used to replace it by the solution. We assume that the vectors c_A and d_A and the values γ_A and δ_A have already been prepared.

computed, we can use them to compute the vectors on level $\ell + 1$, until the maximal level has been reached. The resulting algorithm is given in Figure 3.

```
procedure setup(A);
  if ℓ = 0 then
    Prepare A, e.g., compute its factorization
  else begin
    setup(A_1);
    setup(A_2);
    c_A ← a_2;   solveadj(A_1, c_A)
    d_A ← b_2;   solve(A_2, d_A)
    γ_A ← c_A^* a_1;   δ_A ← γ_A b_1^* d_A
  end
end
```

Figure 3: Setup phase: Prepare the vectors c_A and d_A and the values γ_A and δ_A according to (3) and (5) for all submatrices.

Let us now investigate the complexity of the recursive algorithms. If we denote the storage requirements of the representation (1) of $A \in \mathcal{H}_\ell$ by M_ℓ, we find

$$M_\ell = \begin{cases} n_0^2 & \text{if } \ell = 0, \\ 2M_{\ell-1} + 4n_{\ell-1} = 2M_{\ell-1} + 2n_\ell & \text{otherwise} \end{cases} \qquad \text{for all } \ell \in \mathbb{N}_0,$$

and we can see that this implies

$$M_\ell = (2\ell + n_0)n_\ell \qquad \text{for all } \ell \in \mathbb{N}_0,$$

i.e., if we assume n_0 to be constant, the storage requirements grow like $\mathcal{O}(n_\ell \log n_\ell)$. This is typical for most $\mathcal{H}$-matrix representations.

Lemma 3.1 (Solving) *Assume that there is a constant $C_0 \in \mathbb{R}_{>0}$ such that solving the problems (2) and (7) for level $\ell = 0$ requires not more than $C_0 n_0^2$ operations.*

Then for all $\ell \in \mathbb{N}_0$ and $A \in \mathcal{H}_\ell$, the algorithms given in Figure 1 and Figure 2 require not more than $(C_0 n_0 + 6\ell)n_\ell$ operations.

Proof. We consider only the algorithm given in Figure 1, since both algorithms differ only in the sequence the elementary computation steps are carried out.

We denote the number of operations required on level $\ell \in \mathbb{N}_0$ by $S_\ell \in \mathbb{N}$.

According to our assumption, the algorithm requires not more than $C_0 n_0^2$ operations on level $\ell = 0$, i.e., we have

$$S_0 \leq C_0 n_0^2.$$

Let us now consider a level $\ell > 0$. Computing α_1, α_2 and α_4 each requires $2n_{\ell-1} - 1$ operations, while α_3 is computed in 3 operations, giving us a total of $6n_{\ell-1}$ operations. The updates of x_2 and x_1 each require $2n_{\ell-1}$ operations, giving us $6n_{\ell-1}$ operations for all three updates. Taking the two recursive solves into account, we get

$$S_\ell = 2S_{\ell-1} + 12n_{\ell-1} = 2S_{\ell-1} + 6n_\ell.$$

Now we can use a straightforward induction to prove

$$S_\ell \leq (C_0 n_0 + 6\ell)n_\ell \qquad \text{for all } \ell \in \mathbb{N}_0,$$

and this is the desired estimate. ∎

If we again assume n_0 to be constant, we can see that the number of operations of the solution algorithm grows like $\mathcal{O}(n_\ell \log n_\ell)$, and this can be considered the optimal complexity given that the storage requirements of the matrix show the same asymptotic behaviour.

Lemma 3.2 (Preparing) *Assume that there are constants $C_0, \widehat{C}_0 \in \mathbb{R}_{>0}$ such that solving the problems (2) and (7) for level $\ell = 0$ requires not more than $C_0 n_0^2$ operations and that preparing, e.g., factoring, the matrix A on this level requires not more than $\widehat{C}_0 n_0^3$ operations.*

Then for all $\ell \in \mathbb{N}_0$ and $A \in \mathcal{H}_\ell$, the algorithm given in Figure 3 requires not more than $(\widehat{C}_0 n_0^2 + (C_0 n_0 - 1)\ell + 3\ell^2)n_\ell$ operations.

Proof. We denote the number of operations required on level $\ell \in \mathbb{N}_0$ by $P_\ell \in \mathbb{N}$.

According to our assumption, the algorithm requires not more than $\widehat{C}_0 n_0^3$ operations on level $\ell = 0$, i.e., we have

$$P_0 \leq \widehat{C}_0 n_0^3.$$

Let us now consider a level $\ell > 0$. Due to Lemma 3.1, computing the vectors c_A and d_A takes not more than $(C_0 n_0 + 6(\ell - 1))n_{\ell-1}$ per vector. γ_A is computed using $2n_{\ell-1} - 1$ operations, and δ_A is computed using $2n_{\ell-1}$ operations. This yields

$$\begin{aligned} P_\ell &= 2P_{\ell-1} + 2(C_0 n_0 + 6(\ell-1))n_{\ell-1} + 4n_{\ell-1} - 1 \\ &< 2P_{\ell-1} + (C_0 n_0 + 6\ell - 6)n_\ell + 2n_\ell \\ &= 2P_{\ell-1} + 3(2\ell-1)n_\ell + (C_0 n_0 - 1)n_\ell. \end{aligned}$$

Based on this bound, we can prove

$$P_\ell \leq (\widehat{C}_0 n_0^2 + (C_0 n_0 - 1)\ell + 3\ell^2)n_\ell \qquad \text{for all } \ell \in \mathbb{N}_0$$

by a simple induction, and this is the estimate we need. ∎

Once more assuming that n_0 is constant, the number of operations required to prepare a matrix $A \in \mathcal{H}_\ell$ for the efficient solver grows like $\mathcal{O}(n_\ell \log^2 n_\ell)$. The additional logarithmic factor is introduced since each step of the setup algorithm involves $\mathcal{O}(n_\ell \log n_\ell)$ operations in the solver steps.

Remark 3.3 (Generalization) *The Sherman-Morrison-Woodbury formula can be extended to matrix updates of rank k. In this case the vectors a_1, a_2, b_1, b_2 in (1) can be replaced by matrices of dimension $n_{\ell-1} \times k$, the coefficients γ_A and δ_A become $k \times k$ matrices, and instead of dividing by $1 - \delta_A$, we have to solve a $k \times k$ system, but otherwise the algorithm remains unchanged.*

It is not clear if the algorithm can be extended to more general matrix structures, e.g., those used for three-dimensional integral equations, since this would mean that it is no longer possible to treat the Schur complement by a simple low-rank update.

4 Numerical experiments

Since our algorithm computes the exact solution of the problem (2), we do not have to consider the accuracy of the computed solution, we only have to investigate the runtime behaviour. We consider a simple model problem: A is a symmetric tridiagonal matrix with the value 4 on the diagonal and random values between -1 and 1 on the sub- and superdiagonal. By the Gershgorin circle theorem this guarantees that A is positive definite and therefore $\mathcal{H}$-regular, so our algorithm can be applied.

We use $n_0 = 2$ and consider matrix dimensions up to $n_0 2^{20} = 2097152$. The runtime for preparing the decomposition is shown in Figure 4: the x-axis gives the dimension n_ℓ of the matrix in logarithmic scale, the y-axis gives the time per degree of freedom. We can see that the time grows like $\mathcal{O}(n_\ell \log^2 n_\ell)$, as predicted by our theory.

Figure 5 shows the runtime for solving the linear system once the decomposition has been prepared. We can see that the time grows like $\mathcal{O}(n_\ell \log n_\ell)$, agreeing with our theoretical prediction.

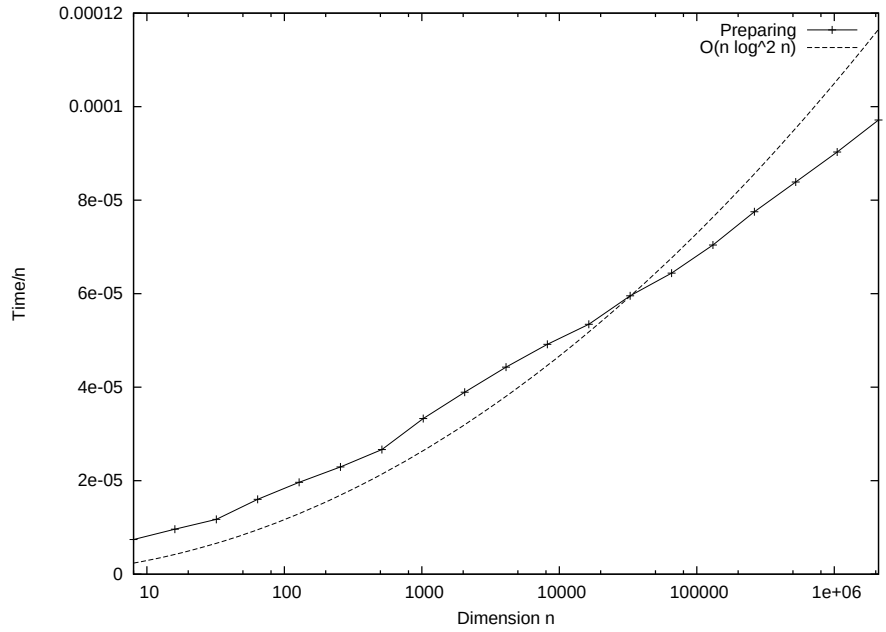

Figure 4: Measured time (per degree of freedom) for preparing the decomposition

References

[1] M. Bebendorf. Why finite element discretizations can be factored by triangular hierarchical matrices. *SIAM J. of Numer. Anal.*, 45(4):1472–1494, 2007.

[2] S. Chandrasekaran, M. Gu, and T. Pals. A fast ULV decomposition solver for hierarchically semiseparable representations. *SIAM J. Matrix Anal. Appl.*, 28(3):603–622, 2006.

[3] L. Grasedyck. Adaptive recompression of $\mathcal{H}$-matrices for BEM. *Computing*, 74(3):205–223, 2004.

[4] L. Grasedyck and W. Hackbusch. Construction and arithmetics of $\mathcal{H}$-matrices. *Computing*, 70:295–334, 2003.

[5] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comp. Phys.*, 73:325–348, 1987.

[6] W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices. *Computing*, 62:89–108, 1999.

[7] W. Hackbusch and B. N. Khoromskij. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part II: Application to multi-dimensional problems. *Computing*, 64:21–47, 2000.

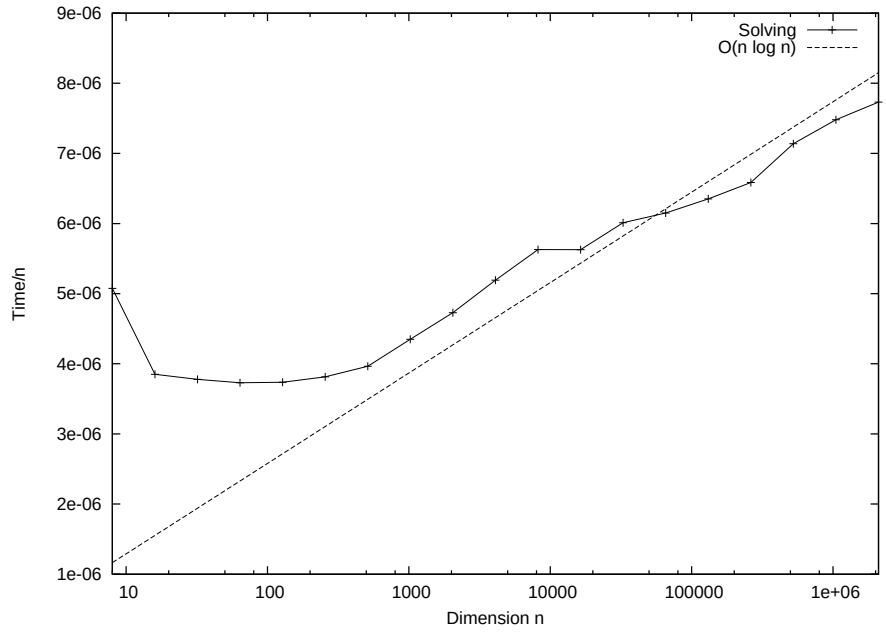

Figure 5: Measured time (per degree of freedom) for solving the linear system

[8] W. Hackbusch and Z. P. Nowak. O cloznosti metoda panelej. In G. I. Marchuk, editor, *Vycislitel'nye prozessy i sistemy*, pages 233–244. Nauka, Moskau, 1988.

[9] W. Hackbusch and Z. P. Nowak. On the fast matrix multiplication in the boundary element method by panel clustering. *Numer. Math.*, 54:463–491, 1989.

[10] S. LeBorne, L. Grasedyck, and R. Kriemann. Domain-decomposition Based H-LU Preconditioners. In O. B. Widlund and D. E. Keyes, editors, *Domain Decomposition Methods in Science and Engineering XVI*, volume 55 of *Lecture Notes in Computational Science and Engineering*, pages 661–668. Springer, 2006.

[11] V. Rokhlin. Rapid solution of integral equations of classical potential theory. *J. Comp. Phys.*, 60:187–207, 1985.

[12] J. Sherman and W. J. Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *Annals of Mathematical Statistics*, 21(1):124–127, 1950.